



DAOS Adventures at CERN openlab

DAOS User Group 2020 (DUG20)

Miguel F. Medeiros *on behalf of openlab team*

miguel.fontes.medeiros@cern.ch

19/11/2020

CERN openlab



- CERN's IT branch for cutting-edge computing technologies.

“CERN openlab is a unique public-private partnership that works to accelerate the development of cutting-edge ICT solutions for the worldwide LHC community and wider scientific research. Through CERN openlab, CERN collaborates with leading ICT companies and research institutes.”

- Active collaboration with **Intel** on several projects: <https://openlab.cern/members/intel>




<https://openlab.cern/>

Disclaimer

- **Work performed under CERN's openlab umbrella**
 - CERN openlab comprises many projects and collaborations which we provide technical support.
 - The specific openlab DAOS use-cases will not be covered on this talk.
- **This talk will only focus on the sysadmin/technical aspects**
 - We will present our experience and process on commissioning, testing and benchmarking DAOS.
 - We will try to provide insights, findings and hopefully valuable feedback for DAOS developers.

Releasing it to our users: the process

- Intel Workshop at CERN on February 2020 (right on time...)
- Benchmark and test it ourselves (with the valuable support of Intel experts)
- **Release DAOS with socket configuration**
 - Allow all users to get acquainted with the system.
 - Test the functionality, development and integration aspects.
- **Release DAOS with PSM2 configuration**
 - Exclusive cluster access.
 - Allow the users with performance requirements to test their use cases.



Cluster Hardware

- **4x Cascade Lakes**

System specifications	
CPU	Intel(R) Xeon(R) Platinum 8260 CPU @ 2.40GHz
CPU per node	24 cores/socket, 2 sockets, 2 threads/core (HT enabled)
Core frequency	Base: 1.0 GHz Range: 1.0GHz - 3.9GHz
Numa nodes	node0: 0-23,48-71 node1: 24-47,72-95
System Memory	12x 32GB DDR4 rank DIMMs
Optane DCPMM	12x 128GB DDR4 rank DIMMs
Optane FW version	01.02.00.5395
BIOS	version: SE5C620.86B.02.01.0011.032620200659 date: 03/26/2020
Storage	4x 1 TB NVMe INTEL SSDPE2KX010T8
HFI	1x Intel Corporation Omni-Path HFI Silicon 100 Series.
HFI Firmware	Termal Management Module: 10.9.0.0.208; Driver: 1.9.2.0.0

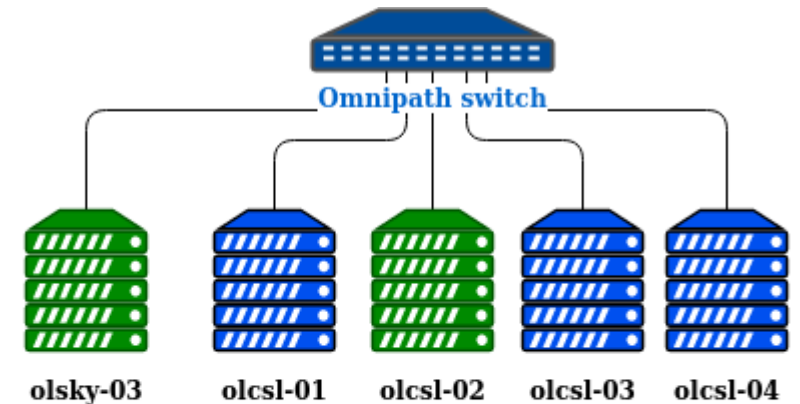
- **4x SkyLakes (only 2x with DAOS)**

System specifications	
CPU	Intel(R) Xeon(R) Platinum 8160 CPU @ 2.10GHz
CPU per node	24 cores/socket, 2 sockets, 2 threads/core (HT enabled)
Core frequency	Base: 1.0 GHz Range: 1.0GHz - 3.9GHz
Numa nodes	node0: 0-23,48-71 node1: 24-47,72-95
System Memory	12x 16GB DDR4 rank DIMMs
BIOS	version: SE5C620.86B.02.01.0011.032620200659 date: 03/26/2020
Storage	1x 350GB INTEL SSDPED1K375GA 1x 480GB INTEL SSDSC2KG480G8
HFI	1x Intel Corporation Omni-Path HFI Silicon 100 Series.
HFI Firmware	Termal Management Module: 10.9.0.0.208; Driver: 1.9.2.0.0



Benchmarking considerations

- Validated the Omnipath cluster with MPI tests & benchmarks
 - OSU Micro-Benchmarks, IntelMPI
- Benchmark was based on IOR [1] with DAOS API
- All benchmarks were performed with **DAOS v0.9.4**
- **Topology**
 - 3x DAOS Servers
 - 2x DAOS Clients


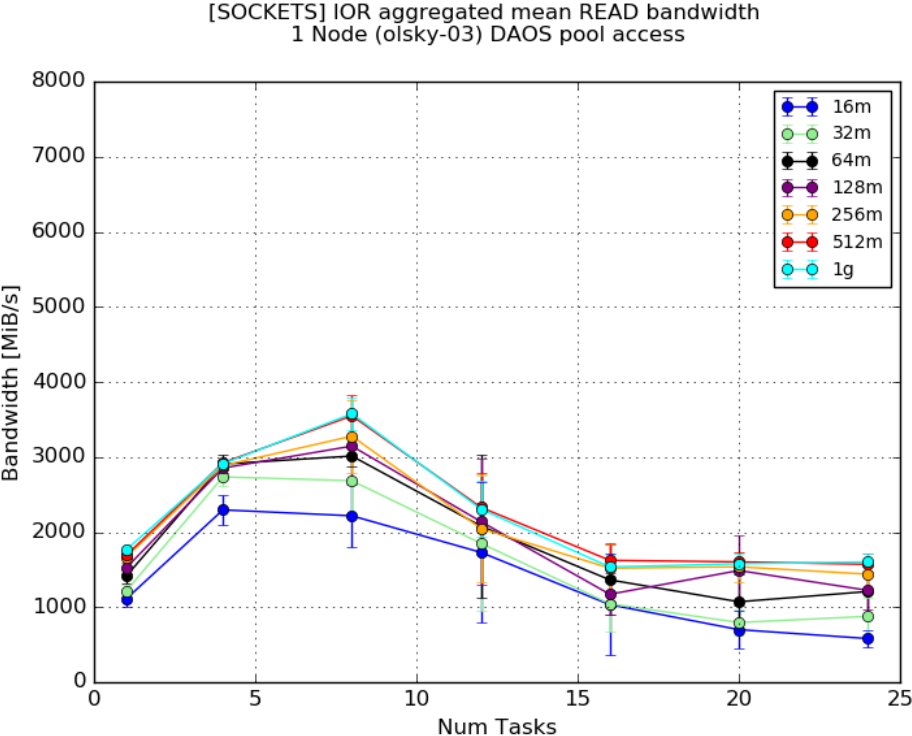


DAOS Agent	Yes	No	Yes	No	No
DAOS Server	No	Yes	No	Yes	Yes
DAOS Head Node	No	No	No	Yes	No

[1]: <https://github.com/hpc/ior>

Benchmarking with IOR: DAOS Sockets

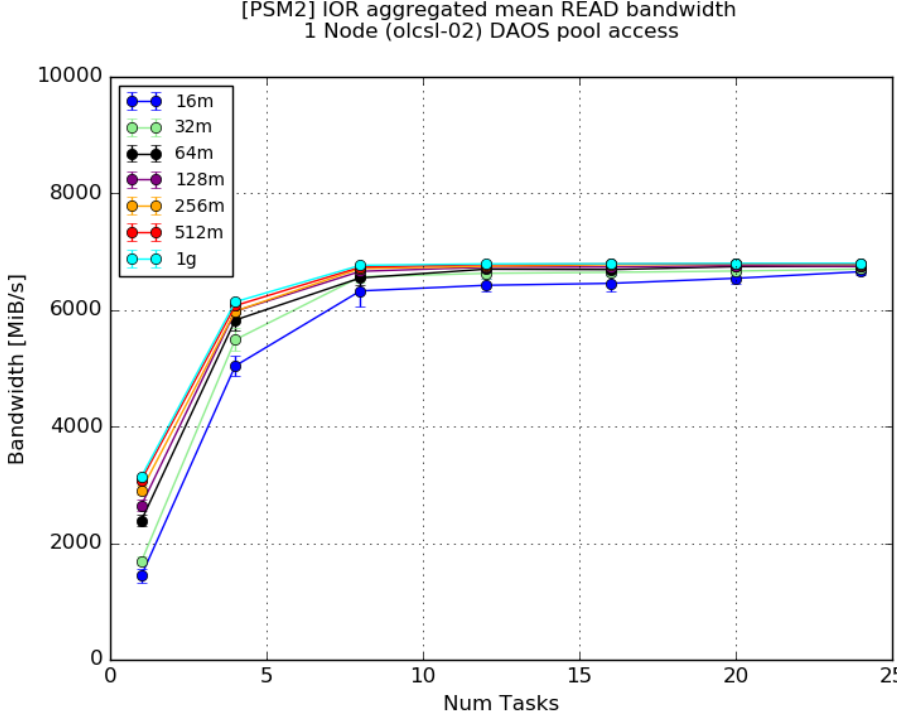
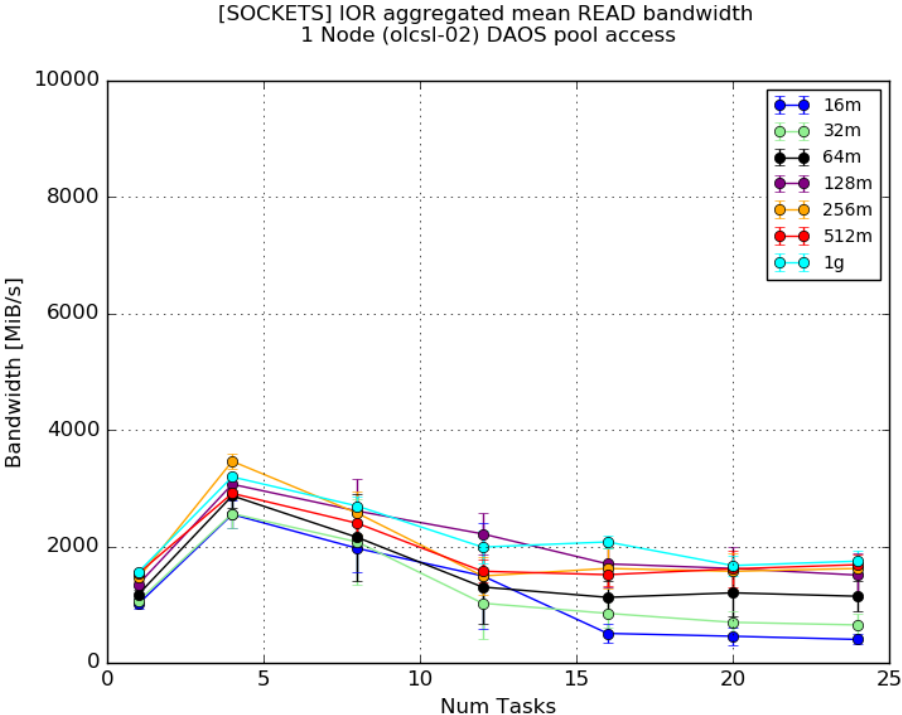
- Functional tests with Sockets configuration



Each point represents the average of 20 iterations. Error bars are standard deviation.

Benchmarking with IOR: DAOS PSM2

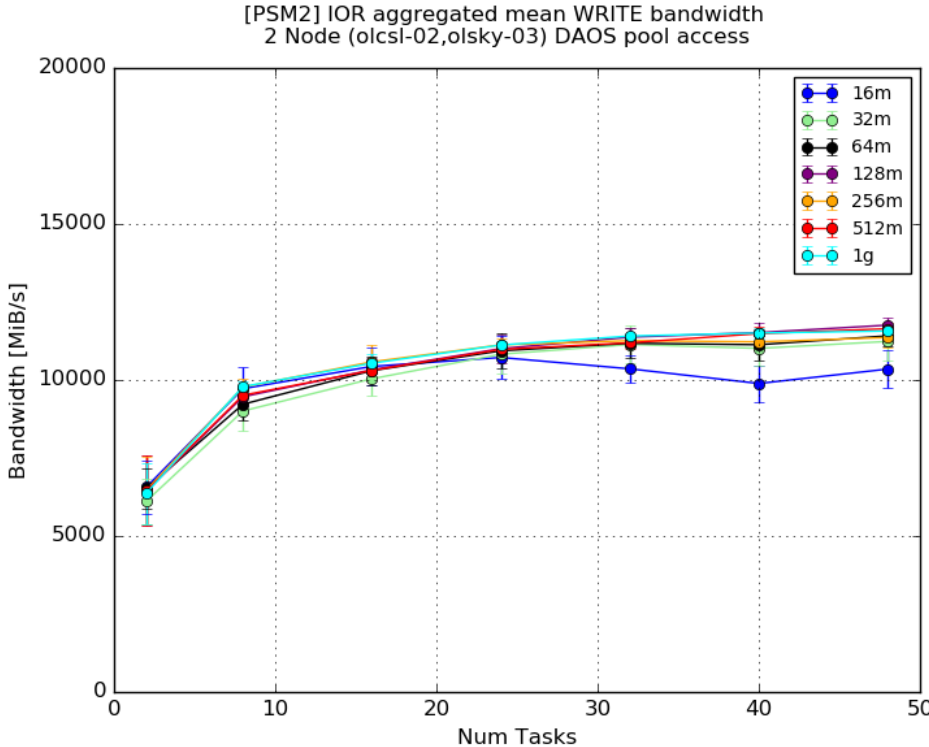
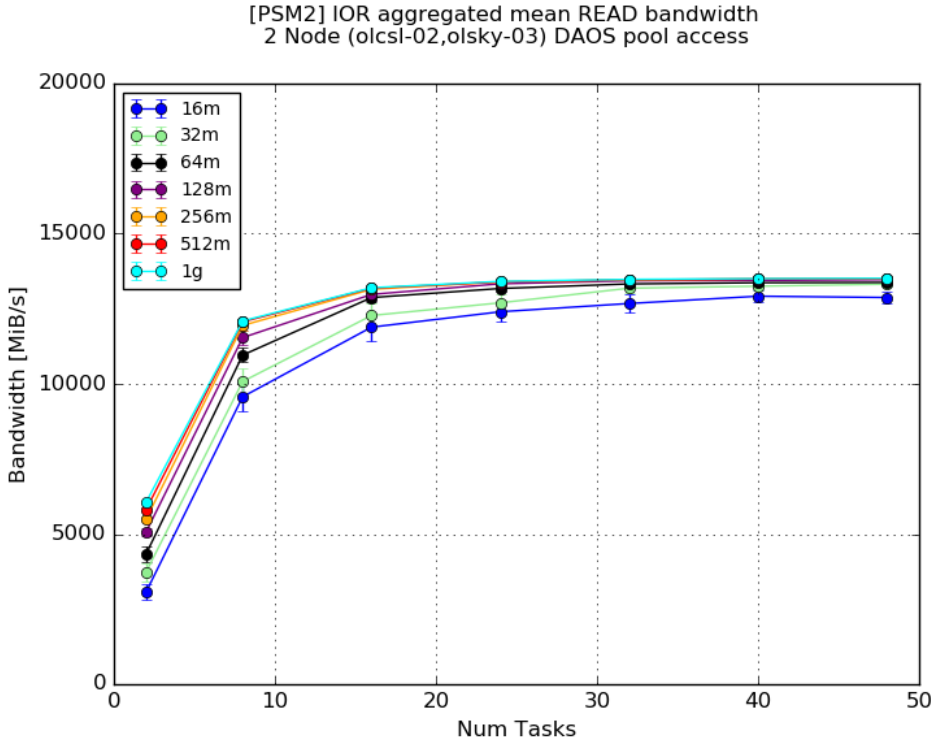
- We also tested with PSM2 → performance gains



Each point represents the average of 20 iterations. Error bars are standard deviation.

Benchmarking with IOR: DAOS PSM2 scaling


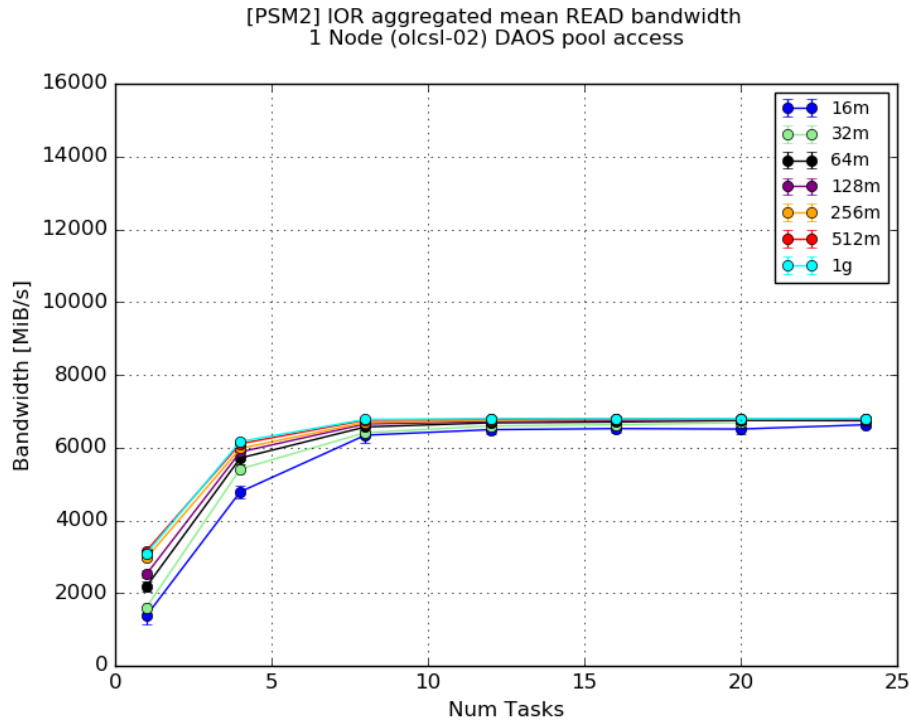
- Two node test



Each point represents the average of 20 iterations. Error bars are standard deviation.

Performance limitations

- Finding the missing performance...
 - Cascade Lake with half performance.
 - We suspect the riser card → HFI card is PCIe 16x but our riser card only provides 8x elec, 8x mech.



Each point represents the average of 20 iterations. Error bars are standard deviation.

A SysAdmin experience – feedback for developers

DISCLAIMER:

- *Please note that all comments provided here are **based on a DAOS v0.9.4 experience**. If something was improved in the meantime please ignore its mention.*

A SysAdmin experience – feedback for developers 1

- **Installation in our environment was challenging**
 - Dependency resolution with conflicts → we rely on specific software versions for our internal software.
 - Warn users about “sanity/pre-flight” checks before compiling.
- **Error reporting**
 - Difficult to troubleshoot some issues.
 - We rely on error messages for troubleshooting and not all errors were mapped on <https://daos-stack.github.io/admin/troubleshooting/#daos-errors>

```
$ dmg -i -l olcsl-03 pool query --pool=c0311e03-e510-4016-a353-15ed028a86e5
olcsl-03:10001: connected
ERROR: dmg: pool query failed: DAOS returned error code: -1005
```

DISCLAIMER: Please note that all comments provided here are based on a DAOS v0.9.4 experience. If something was improved in the meantime please ignore its mention.

A SysAdmin experience – feedback for developers 2

- What about a “*--detail*” option for admins?


```
[mfontesm@olsky-03 ~]$ dmg -i -l olcsl-03 pool list --detail
olcsl-03:10001: connected
Pool UUID                               Svc Replicas  NVMe size  SCM size  Creation date  Creator?...
```

Pool UUID	Svc Replicas	NVMe size	SCM size	Creation date	Creator?...
e6795f86-ff25-4c6e-8877-9a8ebcdd7d78	1				
2366efdf-f22a-444a-8fb6-ab92077cc501	1				
725cc745-557b-46c2-a248-35fef9feda4a	1				
19db26ee-7953-466e-9fde-6c48273f2290	1				
de873392-0926-463f-aaee-5271fe0fa256	1				
5b015554-fdcb-451d-9521-797490864865	1				
003930c7-4ef8-4496-875d-9d884499c02d	1				

DISCLAIMER: Please note that all comments provided here are based on a DAOS v0.9.4 experience. If something was improved in the meantime please ignore its mention.

A SysAdmin experience – feedback for developers 3

- DAOS Servers with constant 30% CPU utilization.



DISCLAIMER: Please note that all comments provided here are based on a DAOS v0.9.4 experience. If something was improved in the meantime please ignore its mention.

A SysAdmin experience – feedback for developers 4

- System metrics
 - Nice to have these metrics in handy formats (e.g: json).
 - Unpack the “message” section.
 - Facilitate the integration on monitoring solutions without extra parsers.

```
[root@olcsl-04 daos]# dmg -i -l olcsl-03 storage query nvme-health --json > test_metric.json
[root@olcsl-04 daos]# jq . test_metric.json
{
  "level": "INFO",
  "time": "2020-11-13T11:59:32+0100",
  "message": "olcsl-03:10001: connected\n"
}
{
  "level": "INFO",
  "time": "2020-11-13T11:59:34+0100",
  "message": "olcsl-03:10001\n\tNVMe controllers and namespaces detail with health statistics:\n\t\tPCI:0000:5e:00.0 Model:INTEL SSDPE2KX010T8 FW:VDV10152 Socket:0 Capacity:1.0 TB\n\t\tHealth Stats:\n\t\t\tTemperature:316K(43C)\n\t\t\tController Busy Time:3h29m0s\n\t\t\tPower Cycles:71\n\t\t\tPower On Duration:9610h0m0s\n\t\t\tUnsafe Shutdowns:57\n\t\t\tMedia Errors:0\n\t\t\tError Log Entries:0\n\t\t\tCritical Warning
s:\n\t\t\tTemperature: OK\n\t\t\tAvailable Spare: OK\n\t\t\tDevice Reliability: OK\n\t\t\tRead Only: OK\n\t\t\tVolatile Memory Backup: OK\n\t\t\tPCI:0000:5f:00.0 Model:INTEL SSDPE2KX010T8 FW:VDV10170 So
cket:0 Capacity:1.0 TB\n\t\tHealth Stats:\n\t\t\tTemperature:316K(43C)\n\t\t\tController Busy Time:1h35m0s\n\t\t\tPower Cycles:17\n\t\t\tPower On Duration:6976h0m0s\n\t\t\tUnsafe Shutdowns:14\n\t\t\tMedia Errors
:0\n\t\t\tError Log Entries:0\n\t\t\tCritical Warnings:\n\t\t\tTemperature: OK\n\t\t\tAvailable Spare: OK\n\t\t\tDevice Reliability: OK\n\t\t\tRead Only: OK\n\t\t\tVolatile Memory Backup: OK\n\t\t\tPCI:0
000:d8:00.0 Model:INTEL SSDPE2KX010T8 FW:VDV10170 Socket:1 Capacity:1.0 TB\n\t\tHealth Stats:\n\t\t\tTemperature:316K(43C)\n\t\t\tController Busy Time:1h31m0s\n\t\t\tPower Cycles:29\n\t\t\tPower On Duration:703
0h0m0s\n\t\t\tUnsafe Shutdowns:25\n\t\t\tMedia Errors:0\n\t\t\tError Log Entries:0\n\t\t\tCritical Warnings:\n\t\t\tTemperature: OK\n\t\t\tAvailable Spare: OK\n\t\t\tDevice Reliability: OK\n\t\t\tRead On
ly: OK\n\t\t\tVolatile Memory Backup: OK\n\t\t\tPCI:0000:d9:00.0 Model:INTEL SSDPE2KX010T8 FW:VDV10170 Socket:1 Capacity:1.0 TB\n\t\tHealth Stats:\n\t\t\tTemperature:315K(42C)\n\t\t\tController Busy Time:1h24m0
s\n\t\t\tPower Cycles:29\n\t\t\tPower On Duration:7030h0m0s\n\t\t\tUnsafe Shutdowns:23\n\t\t\tMedia Errors:0\n\t\t\tError Log Entries:0\n\t\t\tCritical Warnings:\n\t\t\tTemperature: OK\n\t\t\tAvailable Spare
: OK\n\t\t\tDevice Reliability: OK\n\t\t\tRead Only: OK\n\t\t\tVolatile Memory Backup: OK\n"
}
```

DISCLAIMER: Please note that all comments provided here are based on a DAOS v0.9.4 experience. If something was improved in the meantime please ignore its mention.

Final thoughts

- **System commissioning was challenging and interesting!**
 - Required some debugging in our environment.
- **There is still some room for a performance increase**
 - Issue with riser cards, one HFI card per socket, etc.
- **Scalability testing**
 - We needed more nodes to fully evaluate the system.
- **The DAOS server configuration is a plus**
 - Quite user friendly!
 - Works well with configuration management tools (Puppet).
 - The difficulty was mostly to understand which settings suited best for our cluster.
- **On a personal note, it was a good experience with several learning opportunities.**

Thank you Intel for all the support during this Benchmark process!



home.cern